arXiv:1601.01587v1 [cs.DC] 7 Jan 2016

Distributed Synthesis in Continuous Time

Holger Hermanns[1], Jan Krčál[1], and Steen Vester[2]

[1] Saarland University – Computer Science, Saarbrücken, Germany
{hermanns, krcal}@cs.uni-saarland.de
[2] Technical University of Denmark
stve@dtu.dk

Abstract. We introduce a formalism modelling communication of distributed agents strictly in *continuous-time*. Within this framework, we study the problem of synthesising local strategies for individual agents such that a specified set of goal states is reached, or reached with at least a given probability. The flow of time is modelled explicitly based on continuous-time randomness, with two natural implications: First, the non-determinism stemming from interleaving disappears. Second, when we restrict to a subclass of *non-urgent* models, the quantitative value problem for two players can be solved in EXPTIME. Indeed, the explicit continuous time enables players to communicate their states by delaying synchronisation (which is unrestricted for non-urgent models). In general, the problems are undecidable already for two players in the quantitative case and three players in the qualitative case. The qualitative undecidability is shown by a reduction to decentralized POMDPs for which we provide the strongest (and rather surprising) undecidability result so far.

1 Introduction

Distributed self-organising and self-maintaining systems are posing interesting design challenges, and have been subject to many practical [32] as well as theoretical [27, 28] investigations. Distributed systems interact in real time, and one very general way to reason about their timing behaviour is to assume that arbitrary continuous probability distributions govern the timing of local steps as well as of communication steps. We are interested in how foundational properties of such distributed systems differ from models where timing is abstracted. As principal means of communication we consider symmetric handshake communication, since it can embed other forms of communication faithfully [2, 23] including asynchronous and input/output-separated communication.

As an example, consider the problem of leaking a secret from a sandboxed malware to an attacker. The behaviour of attacker and malware (and possibly other components) are prescribed in terms of states, private transitions, labelled synchronisation transitions, and delay transitions which model both local computation times and synchronisation times. The delays are governed by arbitrary continuous probability distributions over real time. Handshake synchronisation is assumed to take place if all devices able to do so agree on the same transition

label. Otherwise the components run fully asynchronously. The sandboxing can be thought of as restricting the set of labels allowed to occur on synchronisation transitions. The question we focus on is how to synthesise the component control strategies for malware and attacker so that they reach their target (of leaking the secret) almost surely or with at least a given probability p.

More precisely, we consider a parallel composition of n modules synchronizing via handshake communication. The modules are modelled by interactive Markov chains (IMCs) [16, 17], a generalization of labelled transition systems and of continuous time Markov chains, equipped with a well-understood compositional theory. Each module may in each state enable private actions, as well as synchronisation actions. It is natural to view such a *distributed IMC* as a game with $n+1$ players, where the last player controls the interleaving of the modules. Each of the other n players controls the decisions in a single module, only based on its local timed history containing only transitions that have occurred within the module. On entering a state of its module, each player selects and commits to executing one of the actions enabled. A private action is executed immediately while a synchronisation action requires a CSP-style handshake [6], it is executed once all modules able to perform this action have committed to it.

For representing delay distributions, we make one decisive and one technical restriction. First, we assume that each distribution is continuous. This for instance disallows deterministic delays of, say, 3 time units. It is an important simplification assumed along our explorations of continuous-time distributed control. Second, we restrict to exponential distributions. This is a pure technicality, since (a) our results can be developed with general continuous distributions, at the price of excessive notational and technical overhead, and (b) exponential distributions can approximate arbitrary continuous distributions arbitrarily close [24]. Together, these assumptions enable us to work in a setting close to interactive Markov chains.

Apart from running in continuous time, the concepts behind distributed IMCs are rather common. Closely related are models based on probabilistic automata [31] or (partially observable) Markov decision processes [3, 25]. In these settings, the power of the interleaving player $n + 1$ is a matter of ongoing debate [7, 8, 26]. The problem is that without additional (and often complicated) assumptions this player is too powerful to be realistic, and can for instance leak information between the other players. This is a problem, e.g. in the security context, making model checking results overly pessimistic [13].

In sharp contrast to the discrete-time settings, in our distributed IMCs the interleaving player $n+1$ does not have decisive influence on the resulting game. The reason is that the interleaving player can only affect the order of transitions between two delays, but neither *which transitions* are taken nor what the different players *observe*. This is rooted in the common alphabet synchronisation and especially the continuous-time nature of the game: the probability of two local modules changing state at the same time instant is zero, except if synchronising.

Example 1. We consider the model displayed on the right where the delay transitions are labelled by some rate λ. It displays a very simplistic malicious *App* trying to communicate a secret to an outside *Att*acker, despite being sandboxed. Innocently looking action *login*, *logout* and *lookup* synchronise *App* and *Att*, while the unlabelled transitions denote some private actions of the respective module.

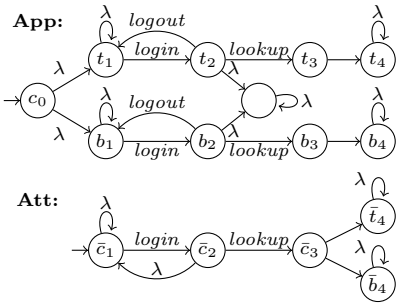

Initially, the *App* can only let time pass. The *Att*acker player has no other choice than committing to handshaking on action *login*. A race of the delay transitions will occur that at some point will lead to either state $(t_1, \bar{c}_1)$ or $(b_1, \bar{c}_1)$, with equal probability. Say in $(t_1, \bar{c}_1)$, the *App* player can only commit to action *login*. The synchronisation will happen immediately since the *Att*acker is committed to *login* already, leading to $(t_2, \bar{c}_2)$. Now the *App* player has either to commit to action *lookup* or *logout*. The latter will induce a deadlock due to a mismatch in players' commitments. Instead assuming the earlier, the state synchronously and immediately changes to $(t_3, \bar{c}_3)$. The *Att*acker player can now use its local timed history to decide which of the private actions to pick. Whatever it chooses, an interleaving of private actions of the two modules follows in zero time. Unless the reachability condition considers transient states such as $(t_3, \bar{t}_4)$ where no time is spent, the player resolving the interleaving has no influence on the outcome.

Now, assume the reachability condition is the state set $\{(t_4, \bar{t}_4), (b_4, \bar{b}_4)\}$. This corresponds to the *Att*acker player correctly determining the initial race of the *App*, and can be considered as a leaked secret. However, according to the explanations provided, it should be obvious that the probability of guessing correctly (by committing properly in state $\bar{c}_3$) is no larger than 0.5, just because the players are bound to decide only based on the local history. The crucial question is: is there an algorithm to compute such probabilities, in general?

Our contribution This paper is the first to explore distributed cooperative reachability games with continuous-time flow modelled explicitly. The formalism we study is based on interactive Markov chains, which in turn has been applied across a wide range of engineering domains. We aim at synthesising local strategies for the players to reach with at least a given probability a specified set of goal states. If this probability is 1 we call the problem *qualitative*, otherwise *quantitative*. We consider *existential* problems, asking for the existence of strategies with these properties, and *value* problems, asking for strategies approximating the given probability value arbitrarily closely. We have three main results:

1. We show that, under mild assumptions on the winning condition, in continuous-time distributed synthesis the interleaving player has no power.
2. In general, we establish that the quantitative problems are undecidable for two or more players, the qualitative value problem is undecidable for two or more players and the qualitative existence problem is EXPTIME-hard for two players and undecidable for three or more players.

3. However, when focusing on the subclass of 2-player *non-urgent* distributed IMCs, the quantitative value problem can be solved in exponential time. Non-urgency enables changing the decisions committed to after some time. Thus, it empowers the players to reach a distributed consensus about the next handshake to perform by observing the only information they jointly have access to: the advance of time.

The qualitative undecidability comes from a novel result about decentralised partially observable Markov decision processes (DEC-POMDP), a multi-player extensions of POMDP. While qualitative existence is decidable for POMDP [1], we show that qualitative existence is undecidable for DEC-POMDP already for 2 players. It is to the knowledge of the authors the strongest undecidability result for DEC-POMDPs with infinite horizon which is of its own interest. By a reduction from DEC-POMDP to distributed IMCs that adds one player, we get undecidability of qualitative existence for 3 or more players in distributed IMCs.

2 Distributed Interactive Markov Chains

We denote by $\mathbb{R}$, $\mathbb{R}_{\geq 0}$, $\mathbb{N}$, and $\mathbb{N}_0$ the sets of real numbers, non-negative real numbers, positive integers, and non-negative integers, respectively. Furthermore, for a finite set X, we denote by $\Delta(X)$ the set of *discrete probability distributions* over X, i.e. functions $f : X \to [0,1]$ such that $\sum_{x \in X} f(x) = 1$. Finally, for a tuple $\boldsymbol{x}$ from a product space $X_1 \times \cdots \times X_n$ and for $1 \leq i \leq n$, we use functional notation $\boldsymbol{x}(i)$ to denote the ith element of the tuple.

We first give a definition of a (local) module based on the formalism of Interactive Markov Chains (IMC). Then we introduce (global) distributed IMC.

Definition 1 (IMC) *An* IMC *(module) is a tuple* $(S, Act, \hookrightarrow, \rightsquigarrow, s^{in})$ *where*

- S *is a finite set of states with an initial state* s^{in}*,*
- Act *is a finite set of actions,*
- $\hookrightarrow \subseteq S \times Act \times S$ *is the* action *transition relation,*
- $\rightsquigarrow \subseteq S \times \mathbb{Q}_{>0} \times S$ *is the finite* delay *transition relation.*

We write $s \overset{a}{\hookrightarrow} s'$ when $(s, a, s') \in \hookrightarrow$ and $s \overset{\lambda}{\rightsquigarrow} s'$ when $(s, \lambda, s') \in \rightsquigarrow$ (λ being the *rate* of the transition). We say that action a is *available* in s if $s \overset{a}{\hookrightarrow} s'$ for some s'.

Definition 2 (Distributed IMC) *A* distributed IMC *is a tuple*

$$\mathcal{G} = ((S_i, Act_i, \hookrightarrow_i, \rightsquigarrow_i, s_i^{in}))_{1 \leq i \leq n}$$

of modules for players $Plr = \{1, ..., n\}$*. Furthermore, by* $Act = \bigcup_i Act_i$ *we denote the set of all actions, and by* $S = S_1 \times ... \times S_n$ *the set of (global) states.*

Intuitively, a distributed IMC moves in continuous-time from a (global) state to a (global) state using transitions with labels from $Label = Act \cup Plr$:

– An action transition with label $a \in Act$ corresponds to *synchronous* communication of all players in $\mathrm{Sync}(a) := \{j \in Plr \mid a \in Act_j\}$ and can only be taken when it is *enabled*, i.e. when all these players choose their local transitions with action a at the same time. It is called a *synchronisation* action if $|\mathrm{Sync}(a)| \geq 2$ and a *private* action if $|\mathrm{Sync}(a)| = 1$.
– A delay transition of any player $j \in Plr$ is taken *independently* by the player after a random delay, i.e. the set of players that synchronise over label j is $\mathrm{Sync}(j) = \{j\}$.

Formally, the (local) *choices of player* j range over $\mathcal{C}_j = \hookrightarrow_j \cup \{\bot\}$. When in (local) state s, the player may pick only a choice *available in* s. That is, either an action transition of the form $s \xhookrightarrow{a} s'$ or $\bot$ if there is no such action transition. We define *global choices* as $\mathcal{C} = \mathcal{C}_1 \times \cdots \times \mathcal{C}_n$. A global choice $\boldsymbol{c}$ induces the set $\mathrm{En}(\boldsymbol{c}) = \{a \in Act \mid \forall j \in \mathrm{Sync}(a) : \boldsymbol{c}(j) = (\cdot, a, \cdot)\}$ of actions *enabled* in $\boldsymbol{c}$.

To avoid that time stops by taking infinitely many action steps in zero time, we pose a standard assumption prohibiting cycles [14,15,18–20]: we require that for every action $a \in Act$ there is a player $j \in \mathrm{Sync}(a)$ such that the labelled transition system $(S_j, \hookrightarrow_j)$ does not have any cycle involving action a.

The behaviour of a distributed IMC is a *play*, an infinite sequence

$$\rho = \boldsymbol{s_0 c_0} \xrightarrow{a_1,t_1} \boldsymbol{s_1 c_1} \xrightarrow{a_2,t_2} \boldsymbol{s_2 c_2} \cdots$$

where each $\boldsymbol{s_i} \in S$ is the state after i moves, $\boldsymbol{c_i} \in \mathcal{C}$ is the choice of the players in the state $\boldsymbol{s_i}$, and $a_{i+1} \in Label$ and $t_{i+1} \in \mathbb{R}_{\geq 0}$ are the label and the absolute time of the next transition taken. By *Play* we denote the set of all plays. Which play is taken depends on the *strategies* of the players, on the *scheduler* which resolves interleaving of communication whenever multiple actions are enabled, and on the rules (involving randomness) given later.

2.1 Schedulers and strategies

First we define strategies and schedulers basing their decision on the current local and global history, respectively. A *(global) history* is a finite prefix

$$h = \boldsymbol{s_0 c_0} \xrightarrow{a_1,t_1} \cdots \xrightarrow{a_i,t_i} \boldsymbol{s_i}$$

of a play ending with a state. For given h, we get the *local history* of player j as

$$\pi_j(h) = \boldsymbol{s'_0}(j)\boldsymbol{c'_0}(j) \xrightarrow{a'_1,t'_1} \cdots \xrightarrow{a'_\ell,t'_\ell} \boldsymbol{s'_\ell}(j)$$

where $\boldsymbol{s'_0 c'_0} \xrightarrow{a'_1,t'_1} \cdots \xrightarrow{a'_\ell,t'_\ell} \boldsymbol{s'_\ell}$ is the subsequence of h omitting all steps not visible for player j, i.e. all $\xrightarrow{a_m,t_m} \boldsymbol{s_m c_m}$ with $j \notin \mathrm{Sync}(a_m)$. The set of all global histories is denoted by *Hist*; the set of local histories of player j by $Hist_j$.

Example 2. Consider again Example 1. Let *App* be controlled by player 1 and *Att* by player 2. For the following history we get corresponding local histories

$$h = (c_0, \bar{c}_1)(\bot, \mathit{login}) \xrightarrow{1,0.42} (t_1, \bar{c}_1)(\mathit{login}, \mathit{login}) \xrightarrow{\mathit{login},0.42} (t_2, \bar{c}_2),$$

$$\pi_1(h) = c_0 \ \bot \xrightarrow{1,0.42} t_1 \ \mathit{login} \xrightarrow{\mathit{login},0.42} t_2, \qquad \pi_2(h) = \bar{c}_1 \ \mathit{login} \xrightarrow{\mathit{login},0.42} \bar{c}_2$$

Note that the attacker can neither observe the Markovian transition nor the local state of the *App*. The *App* cannot observe the local state of the attacker either, but it can be deduced from the local history of the *App*.

A *strategy* for player j is a measurable function $\sigma : \mathit{Hist}_j \to \Delta(\mathcal{C}_j)$ that assigns to any local history h of player j a probability distribution over choices available in the last state of h. We say that a strategy σ for player j is *pure* if for all h we have $\sigma(h)(c) = 1$ for some c; and *memoryless* if for all h and h' with equal last local state we have $\sigma(h) = \sigma(h')$.

A *scheduler* is a measurable function $\delta : \mathit{Hist} \times \mathcal{C} \to \Delta(\mathit{Act}) \cup \{\bot\}$ that assigns to any global history h and global choice $\boldsymbol{c}$ a probability distribution over actions enabled in $\boldsymbol{c}$; or a special symbol $\bot$ again denoting that no action is enabled.

Example 3. The available local choices in $(t_2, \bar{c}_2)$, the last state of h from above, are $\{(t_2, \mathit{lookup}, t_3), (t_2, \mathit{logout}, t_1)\}$ for *App* and solely $\{(\bar{c}_2, \mathit{lookup}, \bar{c}_3)\}$ for *Att*. Let the strategy of *App* select either choice with equal probability. If $(t_2, \mathit{lookup}, t_3)$ is chosen, *lookup* is enabled and must be picked by the scheduler σ. If $(t_2, \mathit{logout}, t_1)$ is chosen, no action is enabled and δ must pick $\bot$, waiting for a delay transition.

2.2 Probability of plays

Let us fix a *profile of strategies* $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ for individual players, and a scheduler δ. The play starts in the initial state $\boldsymbol{s_0} = (s_1^{in}, \ldots, s_n^{in})$ and inductively evolves as follows. Let the current history be $h = \boldsymbol{s_0 c_0} \xrightarrow{a_1, t_1} \cdots \xrightarrow{a_i, t_i} \boldsymbol{s_i}$.

- For the next choice $\boldsymbol{c_i}$, only players $P_i := \mathrm{Sync}(a_i)$ involved in the last transition freely choose (we assume $P_0 := \mathit{Plr}$). Hence, independently for every $j \in P_i$, the choice $\boldsymbol{c_i}(j)$ is taken randomly according to $\sigma_j(\pi_j(h))$. All remaining players $j \notin P_i$ stick to the previous choice $\boldsymbol{c_i}(j) = \boldsymbol{c_{i-1}}(j)$ as for them, no observable event happened.
- After fixing $\boldsymbol{c_i}$, there are two types of transitions:
  1. If $\mathrm{En}(\boldsymbol{c_i}) \neq \emptyset$, the next synchronization action $a_{i+1} \in \mathrm{En}(\boldsymbol{c_i})$ is chosen randomly according to $\delta(h, \boldsymbol{c_i})$ and taken immediately at time $t_{i+1} := t_i$. The next state $\boldsymbol{s_{i+1}}$ satisfies for every $j \in \mathit{Plr}$:

$$\boldsymbol{s_{i+1}}(j) = \begin{cases} \mathrm{target}(\boldsymbol{c_i}(j)) & \text{if } j \in \mathrm{Sync}(a_{i+1}), \\ \boldsymbol{s_i}(j) & \text{if } j \notin \mathrm{Sync}(a_{i+1}). \end{cases}$$

     where $\mathrm{target}(\boldsymbol{c_i}(j))$ denotes the target of the transition chosen by player j. In other words, players involved in synchronisation move according to their choice, the remaining players stay in their previous states.
  2. If $\mathrm{En}(\boldsymbol{c_i}) = \emptyset$, a local delay transition is taken after a random delay, chosen as follows. Each delay transition $\boldsymbol{s_i}(j) \overset{\lambda}{\rightsquigarrow} \cdot$ outgoing from the current local state of any player j is assigned randomly a real-valued delay according to the exponential distribution with rate λ. This results in a collection of real numbers. The transition $\boldsymbol{s_i}(\ell) \overset{\lambda}{\rightsquigarrow} s$ with the minimum

delay d in this collection is taken. Hence, $a_{i+1} := \ell$ (denoting that player ℓ moves), $t_{i+1} := t_i + d$, and the next state $\boldsymbol{s_{i+1}}$ satisfies for every $j \in Plr$:

$$\boldsymbol{s_{i+1}}(j) = \begin{cases} s & \text{if } j \in \mathrm{Sync}(a_{i+1}) = \{\ell\}, \\ \boldsymbol{s_i}(j) & \text{if } j \notin \mathrm{Sync}(a_{i+1}). \end{cases}$$

All these rules induce a probability measure $\Pr^{\boldsymbol{\sigma},\delta}$ over the set of all plays by a standard cylinder construction.

2.3 Distributed synthesis problem

We study the following two fundamental reachability problems for distributed IMCs. Let $\mathcal{G}$ be a distributed IMC, $T \subseteq S$ be a target set of states, and p be a rational number in $[0, 1]$. Denoting by $\diamond T$ the set of plays ρ that reach a state in T and stay there for a non-zero amount of time, we focus on:

Existence Does there exist a strategy profile $\boldsymbol{\sigma}$ s.t. for all schedulers δ,

$$\Pr^{\boldsymbol{\sigma},\delta}(\diamond T) \geq p \ ?$$

Value Can the value p be arbitrarily approached, i.e. do we have

$$\sup_{\boldsymbol{\sigma}} \inf_{\delta} \Pr^{\boldsymbol{\sigma},\delta}(\diamond T) \geq p \ ?$$

We refer to the general problem with $p \in [0, 1]$ as *quantitative*. When we restrict to $p = 1$, we call the problem *qualitative*.

3 Schedulers are not that powerful

The task of a scheduler is to choose among concurrently enabled transitions, thereby resolving the non-determinism conceptually caused by interleaving. In this section, we address the impact of the decisions of the scheduler. We show that despite having the ability to affect the order in which transitions are taken in the *global* play, the scheduler cannot affect what every player observes *locally*. Thus, the scheduler affects neither the choices of any player nor what synchronisation occurs. As a result, for winning objectives that are closed under *local observation equivalence*, the scheduler cannot affect the probability of winning.

Example 4. Consider the distributed IMC to the right. After the delay transition is taken in C_1 and there is synchronisation on action a, the scheduler can choose whether there will be synchronisation on b or c first. However, it can only affect the interleaving, not any of the local plays.

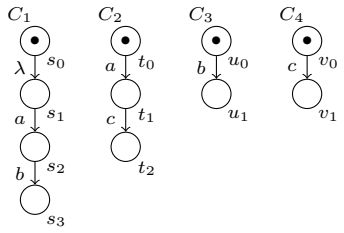

For a play $\rho = \boldsymbol{s_0 c_0} \xrightarrow{a_1,t_1} \boldsymbol{s_1 c_1} \xrightarrow{a_2,t_2} \boldsymbol{s_2 c_2} \cdots$ we define the local play $\pi_j(\rho)$ of player j analogously to local histories. We define *local observation* equivalence $\sim$ over plays by setting $\rho \sim \rho'$ if $\pi_j(\rho) = \pi_j(\rho')$ for all $j \in \mathit{Plr}$. Let us stress that two local observation equivalent plays have exactly the same action and delay transitions happening at the same moments of time; only the order of action transitions happening at the same time can differ. Finally, we say that a set E of plays is *closed under local observation equivalence* if for any $\rho \in E$ and any ρ' such that $\rho \sim \rho'$ we have $\rho' \in E$. It is now possible to show the following.

Theorem 1 *Let E be a measurable set of plays closed under local observation equivalence. For any strategy profile $\boldsymbol{\sigma}$ and schedulers δ and δ' we have*

$$\Pr^{\boldsymbol{\sigma},\delta}(E) = \Pr^{\boldsymbol{\sigma},\delta'}(E).$$

As a result, for the rest of the paper we write $\Pr^{\boldsymbol{\sigma}}(E)$ instead of $\Pr^{\boldsymbol{\sigma},\delta}(E)$ since the scheduler cannot affect the probability of the events we consider. Indeed, the reachability objectives defined in the previous section are closed under local observation equivalence.

Remark 1. The fact that interleaving does not have decisive impact in continuous time may seem natural and thus possibly unsurprising to experts. Yet, the result does not hold for many small variations of the setting we consider, e.g. neither for asymmetric communication nor when allowing cycles of action transitions.

4 Undecidability Results

In this section, we put distributed IMCs into context of other partial-observation models. As a result, we show that reachability quickly gets undecidable here.

Theorem 2 *For distributed IMCs we have that*

1. *the qualitative value, quantitative value, and quantitative existence problems are undecidable with $n \geq 2$ players; and*
2. *the qualitative existence problem is* ExpTime*-hard with $n = 2$ players and undecidable with $n \geq 3$ players.*

Theorem 2 is obtained by using two fundamental results. First, we provide a novel (and somewhat surprising) result for *decentralized POMDPs* (DEC-POMDPs) [3], an established multi-player generalization of POMDPs. We show that the qualitative existence problem for DEC-POMDPs is undecidable already for 2 players. This is, to the knowledge of the authors, currently the strongest known undecidability result for DEC-POMDPs. Second, we show that distributed IMCs are not only more expressive (w.r.t. reachability) than POMDPs but also more expressive than DEC-POMDPs. We show it by reducing reachability in DEC-POMDPs with n players to reachability in distributed IMCs with $n+1$ players. Theorem 2 follows from these two results and from known results about POMDPs [4, 12, 25]. For an overview, see Table 1.

	POMDPs	DEC-POMDPs	Distributed IMCs
Qual. Existence	Dec. [1]	Undec. for ≥ 2 players	Undec. for ≥ 3 players
Value	Undec. [12]	Undec. for ≥ 1 player [12]	Undec. for ≥ 2 players
Quant. Existence	Undec. [25]	Undec. for ≥ 1 player [25]	Undec. for ≥ 2 players
Value	Undec. [4]	Undec. for ≥ 1 player [4]	Undec. for ≥ 2 players

Table 1: Undecidability results for reachability. Unreferenced results are shown here.

4.1 Decentralized POMDP (DEC-POMDP)

We start with a definition of the related formalism of decentralized POMDP [3].

Definition 3 *A DEC-POMDP is a tuple* $(S, Plr, (Act_i, \mathcal{O}_i)_{1 \leq i \leq n}, P, O, s^{in})$ *where*

- S *is a finite set of global states with initial state* $s^{in} \in S$,
- $Plr = \{1, ..., n\}$ *is a finite set of players,*
- Act_i *is a finite set of local actions of player* i *with* $Act_i \cap Act_j = \emptyset$ *if* $j \neq i$, *(by* $Act = Act_1 \times \cdots \times Act_n$ *we denote the set of global actions),*
- $\mathcal{O}_i$ *is a finite set of local* observations *for player* i, *(by* $\mathcal{O} = \mathcal{O}_1 \times \cdots \times \mathcal{O}_n$ *we denote the set of global observations),*
- $P : S \times Act \to \Delta(S)$ *is the transition function which assigns to a state and a global action a probability distribution over successor states, and*
- $O : S \times Act \times S \to \Delta(\mathcal{O})$ *is the* observation function *which assigns to every transition a probability distribution over global observations.*

In contrast to distributed IMCs that capture flow of time explicitly, DEC-POMDP is a discrete-time formalism. A DEC-POMDP starts in the initial state s^{in}. Assuming that the current state is s, one discrete step of the process works as follows. First, each player j chooses an action a_j. Then the next state s' is chosen according to the probability distribution $P(s, \boldsymbol{a})$ where $\boldsymbol{a} = (a_1, \ldots, a_n)$. Then, each player j receives an observation $o_j \in \mathcal{O}_j$ such that the observations $\boldsymbol{o} = (o_1, ..., o_n)$ are chosen with probability $O(s, \boldsymbol{a}, s')(\boldsymbol{o})$. Repeating this forever, we obtain a *play* which is an infinite sequence $\rho = s_0 \boldsymbol{a_0} \boldsymbol{o_0} s_1 \boldsymbol{a_1} \boldsymbol{o_1} \cdots$ where $s_0 = s^{in}$ and for all $i \geq 0$ it holds that $s_i \in S$, $\boldsymbol{a_i} \in Act$, and $\boldsymbol{o_i} \in \mathcal{O}$. Note that the players can only base their decisions on the sequences of observations they receive rather than the actual sequence of states which is not available to them. For a more complete coverage of DEC-POMDPs, see [3].

4.2 Reduction from DEC-POMDP

First we present the reduction from a DEC-POMDP $\mathcal{P}$ to a distributed IMC $\mathcal{G}$. In this subsection, we write $\Pr^{\boldsymbol{\sigma}}_{\mathcal{P}}$ or $\Pr^{\boldsymbol{\sigma}}_{\mathcal{G}}$ instead of $\Pr^{\boldsymbol{\sigma}}$ to distinguish between the probability measure in the DEC-POMDP from the probability measure in the distributed IMC.

Proposition 1 *For a DEC-POMDP $\mathcal{P}$ with n players and a target set T of states of $\mathcal{P}$ we can construct in polynomial time a distributed IMC $\mathcal{G}$ with $n+1$ players and a target set T' of global states in $\mathcal{G}$ where:*

$$\exists \boldsymbol{\sigma} : \Pr^{\boldsymbol{\sigma}}_{\mathcal{G}}(\diamond T) = p \iff \exists \boldsymbol{\sigma}' : \Pr^{\boldsymbol{\sigma}'}_{\mathcal{P}}(\diamond T') = p.$$

Proof (Proof sketch). Let us fix n and $\mathcal{P} = (S, Plr, (Act_i)_{1\leq i\leq n}, \delta, (\mathcal{O}_i)_{1\leq i\leq n}, O)$ where $Plr = \{1, ..., n\}$. Further, let $Act_i = \{a_{i1}, ..., a_{im_i}\}$ and $\mathcal{O}_i = \{o_{i1}, ..., o_{i\ell_i}\}$ for player $i \in Plr$. The distributed IMC $\mathcal{G}$ has $n+1$ modules, one module for each player in $\mathcal{P}$ and the *main* module responsible for their synchronisation. Intuitively,

- the module of every player i stores the last local observation in its state space. Every step of $\mathcal{P}$ is modelled as follows: The player *outputs* to the main module the action it chooses and then *inputs* from the main module the next observation.
- The main module stores the global state in its state space. Every step of $\mathcal{P}$ corresponds to the following: The main module *inputs* the actions of all players one by one, then it randomly picks the new state and new observations according to the rules of $\mathcal{P}$ based on the actions collected. The observations are lastly *output* to all players, again one by one.

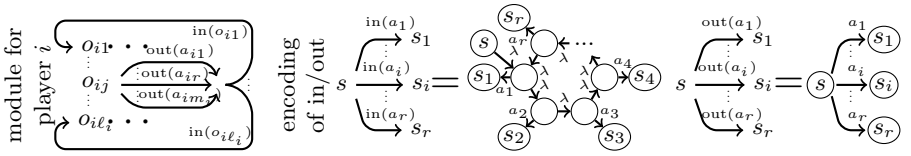

Fig. 1: Module for player i on the left. Input and output encoding to the right.

We construct the distributed IMC so that only the outputting player chooses what action to output whereas the inputting player accepts whatever comes. The construction of modules for player i is illustrated in Figure 1 along with constructions for input and output. The interesting part is how an action from the set $\{a_1, \ldots, a_r\}$ is input in a state s. Instead of waiting in s, the player travels by delay transitions in a round-robin fashion through a cycle of r states, where in the i-th state, only the action a_i is available. Thus, the player has no influence and must input the action that comes. By this construction, the main module has at most one action transition in every state such that the player cannot influence anything; other modules get no insight by observing time and thus the players have the same power as in the DEC-POMDP. □

4.3 Undecidability of qualitative existence in DEC-POMDP

Next, we show that the qualitative existence problem for DEC-POMDPs even with $n \geq 2$ players is undecidable. The proof has similarities with ideas from [5]

where it is shown that deciding existence of sure-winning strategies in safety games with 3 players and partial observation is undecidable. Using the randomness of DEC-POMDPs we show undecidability of the qualitative existence problem for reachability in 2-player DEC-POMDPs.

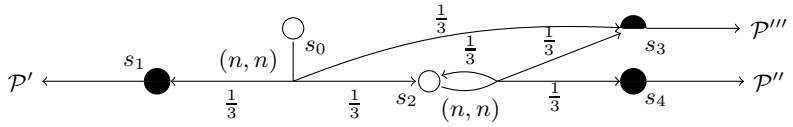

Fig. 2: Overall structure of $\mathcal{P}$ without details of $\mathcal{P}', \mathcal{P}''$ and $\mathcal{P}'''$.

Theorem 3 *It is undecidable whether for a DEC-POMDP $\mathcal{P}$ with $n \geq 2$ players and a set T of target states in $\mathcal{P}$ if there exists a strategy profile $\boldsymbol{\sigma}$ such that* $\Pr^{\boldsymbol{\sigma}}_{\mathcal{P}}(\diamond T) = 1$.

Proof (Proof sketch). We do a reduction from the non-halting problem of a deterministic Turing machine M that starts with a blank input tape. From M we construct a DEC-POMDP $\mathcal{P}$ with two players $Plr = \{1, 2\}$ such that M does not halt if and only if players 1 and 2 have strategies $\boldsymbol{\sigma} = (\sigma_1, \sigma_2)$ which ensure that the probability of reaching a target set T is 1. Figure 2 shows the overall structure of $\mathcal{P}$ without details of sub-modules $\mathcal{P}', \mathcal{P}''$ and $\mathcal{P}'''$.

Both players have two possible observations, black and white. We depict the observation of player 1 in the top-half and of player 2 in the bottom-half of every state. The play starts in s_0 and with probability 1, every player receives the black observation exactly once during the play. If the play goes to s_1 or s_4 the players will receive the observation at the same time and if the play goes to s_3 then player 2 will receive the observation in the step after player 1 does. The modules $\mathcal{P}', \mathcal{P}''$ and $\mathcal{P}'''$ are designed so that:

- In $\mathcal{P}'$, a target state is reached if and only if the sequence of actions played by both players encodes the initial configuration of M.
- In $\mathcal{P}''$, a target state is reached with probability 1 if and only if both players play the same infinite sequence of actions. Note that randomness is essential to build such a module.
- In $\mathcal{P}'''$, the target set is reached if and only if the sequences of actions played by player 1 and 2 encode two configurations C_1 and C_2 of M, respectively, such that C_1 is not an accepting configuration and C_2 is a successor configuration of C_1. This can be done since a finite automaton can be constructed that recognizes if one configuration is a successor of the other when reading both configurations at the same time. Note that it is possible because such configurations can only differ by a constant amount (the control state, the tape head position and in symbols in cells near the tape head).

It can be shown by induction that if there are strategies σ_1, σ_2 that ensure reaching T with probability 1 then every σ_i has to play the encoding of the jth

configuration of M when it receives the black observation in the jth step. Further, it can be shown that these strategies do ensure reaching T with probability 1 if M does not halt on the empty input tape and do not ensure reaching T with probability 1 if M halts. □

5 Decidability for non-urgent models

In this section, we turn our attention to a subclass of distributed IMCs, called *non-urgent*, that implies decidability for both the qualitative and quantitative value problems for 2 players.

Definition 4 *We call* $\mathcal{G} = ((S_i, Act_i, \hookrightarrow_i, \rightsquigarrow_i, s_{0i}))_{1 \le i \le n}$ non-urgent *if for every* $1 \le j \le n$*:*

1. *Every* $s \in S_j$ *is of one of the following forms:*
   (a) Synchronisation *state with at least one outgoing synchronisation action transition and exactly one outgoing delay transition which is a self-loop.*
   (b) Private *state with arbitrary outgoing delay transitions and private action transitions.*
2. *Player* j *has an action* $\varnothing_j \in Act_j$ *enabled in every synchronisation state from* S_j *that allows to "do nothing" and thus postpone the synchronisation. To this end,* $\varnothing_j$ *is also in* Act_k *for every other player* $k \neq j$ *but* $\varnothing_j$ *does not appear in* $\hookrightarrow_k$*. As a result,* j *does not take part in any synchronisation while choosing* $\varnothing_j$*.*

In a non-urgent distributed IMC, $\boldsymbol{s} \in S$ is called a (global) *synchronisation* state if it is the initial state or all $\boldsymbol{s}(j)$ are synchronisation states. We denote global synchronisation states by S'. All other global states $S \setminus S'$ are called *private*.

Example 5. Consider the non-urgent variant of Example 1 on the right. The "do nothing" actions are a natural concept; the only real modelling restriction is that one cannot model a communication time-out any more, the delay transitions from synchronisation states need to be self-loops.

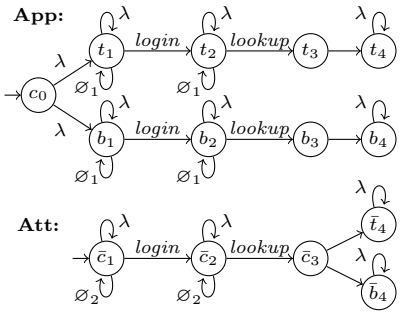

Surprisingly, in this model, the secret can be leaked with probability 1 as follows. As before, the players reach the states $(t_2, \bar{c}_2)$ or $(b_2, \bar{c}_2)$ with equal probability. Now, the *App* player can arbitrarily postpone the *lookup* by committing to action $\varnothing_1$. Whenever the delay self-loop is taken, the player can re-decide to perform *lookup*. Since the self-loop is taken repetitively, the *App* player is flexible in choosing the timing of *lookup*. Thus, leaking the secret is simple, e.g. by performing *lookup* in an odd second when in t_2 and in an even second when in b_2.

For two players, we construct a general synchronisation scheme that (highly probably) shows the players the current global state after each communication.

Theorem 4 *The quantitative value problem for 2-player non-urgent distributed IMCs where the target set consists only of synchronisation states is in* ExpTime.

Being a special case, also the qualitative value problem is decidable. In essence, the problem becomes decidable because in the synchronisation states, the players can effectively exchange arbitrary information. This resembles the setting of [11]. The insight that observing global time provides an additional synchronization mechanism is not novel in itself, but it is obviously burdensome to formally capture in time-abstract models of asynchronous communication, and thus usually not considered. For distributed IMC, it still is non-trivial to develop; here it hinges on the non-urgency assumption. The results of [11] also indicate that for three or more players, additional constraints on the topology may be needed to obtain decidability.

In the rest of the section we prove Theorem 4, fixing a 2-player non-urgent distributed IMC $\mathcal{G} = ((S_i, Act_i, \hookrightarrow_i, \rightsquigarrow_i, s_{0i}))_{1 \leq i \leq 2}$, $p \in [0,1]$, and $T \subseteq S'$. We present the algorithm based on a reduction to a discrete-time Markov decision process and then discuss its correctness.

Markov decision process (MDP) An MDP is a tuple $\mathcal{M} = (S, A, P, s_0)$ where S is a finite set of states, A is a finite set of actions, $P : S \times A \to \Delta(S)$ is a partial probabilistic transition function, and s_0 is an initial state. An MDP is the special case of a DEC-POMDP with 1 player that has a unique observation for each state. A *play* in $\mathcal{M}$ is a sequence $\omega = s_0 s_1 \ldots$ of states such that $P(s_i, a_i)(s_{i+1}) > 0$ for some action a_i for every $i \geq 0$. A history is a prefix of a play. A *strategy* is a function π that to every history $h \cdot s$ assigns a probability distribution over actions such that if an action a is assigned a non-zero probability, then $P(s, a)$ is defined. A strategy π is *pure memoryless* if it assigns Dirac distributions to any history and its choice depends only on the last state of the history. When we fix a strategy π, we obtain a probability measure $\Pr^\pi$ over the set of plays. For further details, see [29].

The algorithm It works by reduction to an MDP $\mathcal{M}_\mathcal{G} = (S', A, P, \boldsymbol{s_0})$ where

- $S' \subseteq S$ is the set of global synchronisation states;
- $A = \mathcal{C} \times \Sigma_1 \times \Sigma_2 \cup \{\bot\}$ where Σ_j is the set of pure memoryless strategies of player j in $\mathcal{G}$ that choose $\varnothing_j$ in every synchronisation state;
- For an arbitrary state (s_1, s_2), we define the transition function as follows:
  - For any $(\boldsymbol{c}, \sigma_1, \sigma_2) \in A$, the transition $P((s_1, s_2), (\boldsymbol{c}, \sigma_1, \sigma_2))$ is defined if $\boldsymbol{c}$ is available in (s_1, s_2) and the players agree in $\boldsymbol{c}$ on some action a, i.e. $\mathrm{En}(\boldsymbol{c}) = \{a\}$. If defined, the distribution $P((s_1, s_2), (\boldsymbol{c}, \sigma_1, \sigma_2))$ assigns to any successor state (s_1', s_2') the probability that in $\mathcal{G}$ the state (s_1', s_2') is reached from (s_1, s_2) via states in $S \setminus S'$ by choosing $\boldsymbol{c}$ and then using the pure memoryless strategy profile (σ_1, σ_2).
  - To avoid deadlocks, the transition $P((s_1, s_2), \bot)$ is defined iff no other transition is defined in (s_1, s_2) and it is a self-loop, i.e. it assigns probability 1 to (s_1, s_2).

The MDP $\mathcal{M}_{\mathcal{G}}$ has size exponential in $|\mathcal{G}|$. Note that all target states T are included in S'. Slightly abusing notation, let $\diamond T$ denote the set of plays in $\mathcal{M}_{\mathcal{G}}$ that reach the set T. From standard results on MDPs [29], there exists an optimal pure memoryless strategy π^*, i.e. a strategy satisfying $\Pr^{\pi^*}(\diamond T) = \sup_{\pi} \Pr^{\pi}(\diamond T)$. Furthermore, such a strategy π^* and the value $v := \Pr^{\pi^*}(\diamond T)$ can be computed in time polynomial in $|\mathcal{M}_{\mathcal{G}}|$. Finally, the algorithm returns TRUE if $v \geq p$ and FALSE, otherwise.

Correctness of the algorithm Let us explain why the approach is correct.

Proposition 2 *The value of $\mathcal{G}$ is equal to the value of $\mathcal{M}_{\mathcal{G}}$, i.e.*

$$\sup_{\boldsymbol{\sigma}} \Pr^{\boldsymbol{\sigma}}(\diamond T) \;=\; \sup_{\pi} \Pr^{\pi}(\diamond T).$$

Proof sketch. As regards the $\leq$ inequality, it suffices to show that any strategy profile $\boldsymbol{\sigma}$ can be mimicked by some strategy π. This is simple as π in $\mathcal{M}_{\mathcal{G}}$ has always knowledge of the global state. Much more interesting is the $\geq$ inequality. We argue that for any strategy π there is a sequence of local strategy profiles $\boldsymbol{\sigma}^1, \boldsymbol{\sigma}^2, \ldots$ such that

$$\lim_{i \to \infty} \Pr^{\boldsymbol{\sigma}^i}(\diamond T) = \Pr^{\pi}(\diamond T).$$

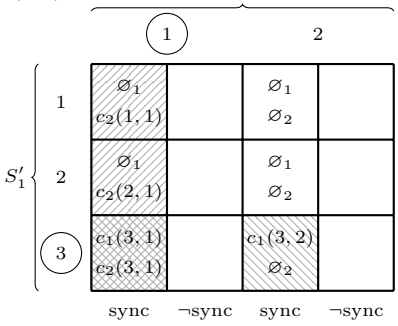

The crucial idea is that a strategy profile communicates correctly (with high probability) the current global state in a synchronisation state by *delaying* as follows. The time is divided into phases, each of $|S'_1| \cdot 2|S'_2|$ slots (where S'_i are the synchronisation states of player i). We depict a phase by the table on the right where the time flows from top to bottom and from left to right (as reading a book). Players 1 and 2 try to synchronise in the row and column, respectively, corresponding to their current states (in circle) and in each slot take the choice $c_i(s_1, s_2)$ optimal given the current global state is (s_1, s_2); in the remaining slots they choose to do nothing. Since the players can change their choice only at random moments of time, their synchronising choice always stretches a bit into the successive silent slot (in a ¬sync column). The more we increase the size of each slot, the lower is the chance that a synchronisation choice of a player stretches to the next *synchronisation* slot. Thus, the lower is the chance of an erroneous synchronisation. We define the size of the slot to increase with i and also along the play so that for any fixed i the probability of at least one erroneous synchronisation is bounded by $\kappa_i < 1$ and for $i \to \infty$, we have $\kappa_i \to 0$. □

6 Discussion and Conclusion

This paper has introduced a foundational framework for modelling and synthesising distributed controllers interacting in continuous time via handshake

communication. The continuous time nature of the model induces that the interleaving scheduler has in fact very little power. We studied cooperative reachability problems for which we presented a number of undecidability results, while for non-urgent models we established decidability of both quantitative and qualitative problems for the two-player case. In the framework considered, the restriction to exponential distributions is a technical vehicle, the results could have been developed in general continuous time, e.g. by using the model of stochastic automata [9].

Distributed IMCs can be considered as an attractive base model especially in the context of information flow and other security-related studies. This is because in contrast to the discrete time setting, the power of the interleaving scheduler is no matter of debate, it can leak no essential information.

From a more general perspective, distributed synthesis of control algorithms has received considerable attention in its entirety [21, 27, 28]. The asynchronous setting with handshake synchronisation has been considered in [22]. Notably, our assumption that players stay committed to a particular action choice for the time in between state changes implies the necessity to let the players explicitly solve distributed consensus problems. As done in [22], one can overcome this by letting local players pick sets of enabled actions (or letting them change choice with infinite speed), and then let some built-in magic pick a valid action from the intersection, implying that whenever possible a consensus is reached for sure. Such a change would however reintroduce the scheduler.

We should point out that distributed IMCs are not fully compositional: We are assuming a fixed vector of modules, and do not discuss that modules themselves may be vectors. Otherwise, we would face the phenomenon of auto-concurrency [10], where transitions with identical synchronisation actions might get enabled concurrently, despite not synchronising. This in turn would again re-introduce distinguishing power of the scheduler.

Distributed Markov chains [30] constitute another recent discrete-time approach where interleaving nondeterminism is tamed successfully via assumptions on the communication structure. The observation that continuous time reduces the power of the interleaving scheduler is not entirely new. Though not explicitly discussed, it underpins the model of probabilistic I/O automata (PIOA) [33] which uses I/O communication with input-enabledness and output-determinism. In that setting, output-determinism implies that local players have no decisive power, and hence a continuous time Markov chain arises. We can approximate I/O-based communication by distributed IMCs without the need for output-determinism. The approximation is linked to arbitrarily small but non-zero delays needed to cycle through synchronising action sets. A profound investigation of the continuous-time particularities of this and other synchronisation disciplines is considered an interesting topic for future work.

Acknowledgements This work is supported by the EU 7th Framework Programme projects 295261 (MEALS) and 318490 (SENSATION), by the Czech Science Foundation project P202/12/G061, the DFG Transregional Collaborative Research Centre SFB/TR 14 AVACS, and by the CDZ project 1023 (CAP).

References

1. Christel Baier, Marcus Größer, and Nathalie Bertrand. Probabilistic ω-automata. *Journal of the ACM*, 59(1):1, 2012.
2. Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
3. Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4):819–840, 2002.
4. Alberto Bertoni, Giancarlo Mauri, and Mauro Torelli. Some recursive unsolvable problems relating to isolated cutpoints in probabilistic automata. In *ICALP*, volume 1256 of *LNCS*, pages 87–94. Springer, 1977.
5. Dietmar Berwanger and Lukasz Kaiser. Information tracking in games on graphs. *Journal of Logic, Language and Information*, 19(4):395–412, 2010.
6. Stephen D. Brookes, C. A. R. Hoare, and A. W. Roscoe. A theory of communicating sequential processes. *Journal of the ACM*, 31(3):560–599, 1984.
7. Ran Canetti, Ling Cheung, Dilsun Kirli Kaynar, Moses Liskov, Nancy A. Lynch, Olivier Pereira, and Roberto Segala. Analyzing security protocols using time-bounded task-PIOAs. *Discrete Event Dynamic Systems*, 18(1):111–159, 2008.
8. Ling Cheung, Nancy A. Lynch, Roberto Segala, and Frits W. Vaandrager. Switched PIOA: parallel composition via distributed scheduling. *Theoretical Computer Science*, 365(1-2):83–108, 2006.
9. Pedro R. D'Argenio and Joost-Pieter Katoen. A theory of stochastic systems part I: stochastic automata. *Information and Computation*, 203(1):1–38, 2005.
10. Manfred Droste and Paul Gastin. Asynchronous cellular automata for pomsets without auto-concurrency. In *CONCUR*, volume 1119 of *LNCS*, pages 627–638. Springer, 1996.
11. Blaise Genest, Hugo Gimbert, Anca Muscholl, and Igor Walukiewicz. Asynchronous games over tree architectures. In *ICALP*, volume 7966 of *LNCS*, pages 275–286. Springer, 2013.
12. Hugo Gimbert and Youssouf Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP*, volume 6199 of *LNCS*, pages 527–538. Springer, 2010.
13. Sergio Giro, Pedro R. D'Argenio, and Luis María Ferrer Fioriti. Distributed probabilistic input/output automata: Expressiveness, (un)decidability and algorithms. *Theoretical Computer Science*, 538:84–102, 2014.
14. D. Guck, T. Han, J.-P. Katoen, , and M.R. Neuhäußer. Quantitative timed analysis of interactive markov chains. In *NFM*, volume 7226 of *LNCS*, pages 8–23. Springer, 2012.
15. H. Hermanns and S. Johr. May we reach it? Or must we? In what time? With what probability? In *MMB*, pages 125–140. VDE Verlag, 2008.
16. Holger Hermanns. *Interactive Markov Chains And The Quest for Quantified Quality*, volume 2428 of *LNCS*. Springer, 2002.
17. Holger Hermanns and Joost-Pieter Katoen. The How and Why of interactive Markov chains. In *Formal Methods for Components and Objects*, volume 6286 of *LNCS*, pages 311–337. Springer, 2009.
18. Holger Hermanns, Jan Krcál, and Jan Kretínský. Compositional verification and optimization of interactive Markov chains. In *CONCUR*, volume 8052 of *LNCS*, pages 364–379. Springer, 2014.

19. J.-P. Katoen, D. Klink, and M. R. Neuhäußer. Compositional abstraction for stochastic systems. In *FORMATS*, volume 5813 of *LNCS*, pages 195–211. Springer, 2009.
20. J.-P. Katoen, I.S. Zapreev, E.M. Hahn, H. Hermanns, and D.N. Jansen. The ins and outs of the probabilistic model checker MRMC. *Performance Evaluation*, 68(2):90–104, 2011.
21. P. Madhusudan and P. S. Thiagarajan. Distributed controller synthesis for local specifications. In *ICALP*, volume 2076 of *LNCS*, pages 396–407. Springer, 2001.
22. P. Madhusudan and P. S. Thiagarajan. A decidable class of asynchronous distributed controllers. In *CONCUR*, volume 2421 of *LNCS*, pages 145–160. Springer, 2002.
23. Robin Milner. Calculi for synchrony and asynchrony. *Theoretical Computer Science*, 25:267–310, 1983.
24. Marcel F Neuts. *Matrix-geometric solutions in stochastic models: an algorithmic approach*. Courier Corporation, 1981.
25. Azaria Paz. *Introduction to Probabilistic Automata*. Academic Press, Inc., 1971.
26. Silvia S. Pelozo and Pedro R. D'Argenio. Security analysis in probabilistic distributed protocols via bounded reachability. In *TGC*, volume 8191 of *LNCS*, pages 182–197. Springer, 2012.
27. Amir Pnueli and Roni Rosner. On the synthesis of an asynchronous reactive module. In *ICALP*, volume 372 of *LNCS*, pages 652–671. Springer, 1989.
28. Amir Pnueli and Roni Rosner. Distributed reactive systems are hard to synthesize. In *FOCS*, pages 746–757. IEEE Computer Society, 1990.
29. Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*, volume 414. John Wiley & Sons, 2009.
30. Ratul Saha, Javier Esparza, Sumit Kumar Jha, Madhavan Mukund, and P. S. Thiagarajan. Distributed Markov chains. In *VMCAI*, volume 8931 of *LNCS*, pages 117–134. Springer, 2015.
31. Roberto Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Massachusetts Institute of Technology, 1995.
32. Bruno Sinopoli, Courtney Sharp, Luca Schenato, Shawn Schaffert, and S. Shankar Sastry. Distributed control applications within sensor networks. *Proc. of the IEEE*, 91:1235–1246, 2003.
33. Sue-Hwey Wu, Scott A. Smolka, and Eugene W. Stark. Composition and behaviors of probabilistic I/O automata. *Theoretical Computer Science*, 176(1-2):1–38, 1997.

A Definition of the probability measure

First, we pose an assumption on the module of each player j to avoid that a play stops because no further steps are possible.

Assumption 1 *We assume that every state has at least one outgoing delay transition or it only has outgoing private action transitions (that cannot be blocked by other modules).*

This is no real restriction: for states without any transition, we can add a delay self-loop; states with synchronisation transitions are transformed as depicted in Figure 3. Note that in the second case, simply adding delay self-loops would change the behaviour because taking a delay self-loop allows the player to change the choice.

As common for continuous-time systems, the definition is based on *cylinder* sets generated from interval-timed histories. An *interval-timed history* is a sequence

$$H = \boldsymbol{s_0 c_0} \xrightarrow{a_1, I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k, I_k} \boldsymbol{s_k}$$

where each I_{i+1} is a real interval bounding the time spent *waiting in* $\boldsymbol{s_i}$. We further require that $I_i = [0, 0]$ whenever $a_i \in \mathit{Act}$ and that the set of histories that *conform* to H is non-empty. We say that a history $\boldsymbol{s_0 c_0} \xrightarrow{a_1, t_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k, t_k} \boldsymbol{s_k}$ *conforms to* H if $t_i - t_{i-1} \in I_i$ for each $1 \leq i \leq k$ (where $t_0 = 0$).

Slightly abusing notation, we interpret an interval-timed history H also as the set of histories that conform to H. We define the cylinder $\mathrm{Cyl}(H) = \{\rho \in \mathit{Play} \mid \exists i.\rho_{\leq i} \in H\}$ as the set of plays which have a prefix in H. We define the measurable spaces $(\mathit{Play}, \mathcal{F})$ and $(\mathit{Hist}, \mathcal{H})$ where $\mathcal{F}$ and $\mathcal{H}$ are the σ-algebra generated from all cylinders and interval-timed histories, respectively:

$$\mathcal{F} = \sigma\left(\{\mathrm{Cyl}(H) \mid H \text{ is an interval-timed history}\}\right)$$
$$\mathcal{H} = \sigma\left(\{H \mid H \text{ is an interval-timed history}\}\right)$$

Analogously, we can define measurable spaces $(\mathit{Hist}_j, \mathcal{H}_j)$ over local histories (by allowing also I_i for action transitions to have non-zero length).

For a given *strategy profile* $\boldsymbol{\sigma}$, i.e. a tuple of strategies $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ of individual players, a given scheduler δ and initial state $\boldsymbol{s_0}$, we obtain a purely

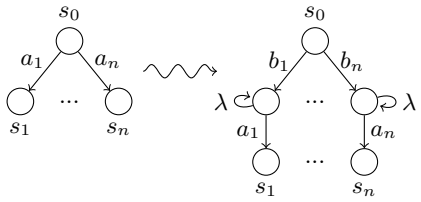

Fig. 3: Transformation for states with synchronisation transitions that adds new private actions $b_1, \ldots, b_n$.

stochastic process and we can define a probability measure $\Pr_{\boldsymbol{s_0}}^{\boldsymbol{\sigma},\delta}$ over *Play*. The probability measure is uniquely determined by fixing probabilities for cylinder sets $\mathrm{Cyl}(H)$ for any interval-timed history

$$H = \boldsymbol{s_0 c_0} \xrightarrow{a_1,I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k,I_k} \boldsymbol{s_k}.$$

Let $\alpha_1, ..., \alpha_v$ be the indices of delay transitions, i.e. for each α_j we have $a_{\alpha_j} \in \mathit{Plr}$; and $\beta_1, ..., \beta_u$ be the indices of action transitions, i.e. for each β_j we have $a_{\beta_j} \in \mathit{Act}$. For $1 \leq i \leq v$, let $\ell_i = \inf(I_{\alpha_i})$ and $u_i = \sup(I_{\alpha_i})$. Then $\Pr_{\boldsymbol{s_0}}^{\boldsymbol{\sigma},\delta}(\mathrm{Cyl}(H))$ is defined as

$$\int_{\ell_1}^{u_1} ... \int_{\ell_v}^{u_v} \prod_{1 \leq i \leq v} De^{\alpha_i}(d_i) \cdot \prod_{0 \leq i < k} St^i \cdot \prod_{1 \leq i \leq u} Sc^{\beta_i} \, \mathrm{d}d_v \cdots \mathrm{d}d_1$$

where the terms $De^i(d)$, Sc^i, and St^i express the contribution of the ith transition to the overall probability caused by the delays, decisions of the scheduler, and decision on the strategies, respectively (see below). The variables d_i denote the delay at ith delay transitions and induce a history

$$h = \boldsymbol{s_0 c_0} \xrightarrow{a_1,t_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k,t_k} \boldsymbol{s_k}$$

such that $t_i = \sum_{\ell, \alpha_\ell \leq i} d_{\alpha_\ell}$. Finally, we set

$$De^i(d) = Q(\boldsymbol{s}_i, \boldsymbol{s}_{i+1}) \cdot e^{-E(\boldsymbol{s}_i)d},$$

where $Q(\boldsymbol{s}, \boldsymbol{s'}) = \sum_{\boldsymbol{s} \overset{\lambda,j}{\rightsquigarrow} \boldsymbol{s'}} \lambda$ and $E(\boldsymbol{s}) = \sum_{\boldsymbol{s'} \neq \boldsymbol{s}} Q(\boldsymbol{s}, \boldsymbol{s'})$, and

$$St^i = \prod_{j \in \mathrm{Sync}(a_i)} \sigma_j(\pi_j(h_{\leq i}))(\boldsymbol{c_i}(j)),$$
$$Sc^i = \delta(h_{\leq i-1}, \boldsymbol{c_{i-1}})(a_i).$$

When $\boldsymbol{s'_0} \neq \boldsymbol{s_0}$ then $\Pr_{\boldsymbol{s'_0}}^{\boldsymbol{\sigma},\delta}(\mathrm{Cyl}_{\mathcal{G}}(H)) = 0$.

B Proof of Theorem 1

Before we prove Theorem 1 we need a few definitions and lemmas.

For an interval-timed history let $\pi_j(H) = \{\pi_j(\rho) \in \mathit{Play} \mid \rho \text{ conforms to } H\}$. Further, Let $\sim$ be an equivalence relation on interval-timed histories defined such that $H \sim H'$ for two interval-timed histories H and H' if and only if $\pi_j(H) = \pi_j(H')$ for every $j \in \mathit{Plr}$. We write $[H]$ for the set of histories h such that there exists $H' \sim H$ with $h \in H'$.

Now, for an interval-timed history $H = \boldsymbol{s_0 c_0} \xrightarrow{a_1,I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k,I_k} \boldsymbol{s_k}$ such that the last action $a_k \in \mathit{Plr}$ if $k > 0$ we define the interleaving-abstract cylinder as the set of plays with a prefix conforming to interval-timed histories H' assuring $H \sim H'$:

$$\mathrm{Cyl}^{\mathrm{ia}}(H) = \{\rho \in \mathit{Play} \mid \rho_{\leq k} \in [H]\}$$

Note that the interleaving-abstract cylinders are contained in the σ-algebra $\mathcal{F}$ generated by the cylinder sets. We can therefore define a sub-σ-algebra $\mathcal{I}$ of $\mathcal{F}$ generated by interleaving-abstract cylinders

$$\mathcal{I} = \sigma(\{\mathrm{Cyl}^{\mathrm{ia}}(H) \mid H \text{ is an interval-timed history s.t. the last action is not in } Act\})$$

Since this is a sub-σ-algebra of $\mathcal{F}$ it inherits the probability measures defined earlier restricted to $\mathcal{I}$. Note that interleaving-abstract events are also 0-time abstract. That is, they are events which are invariant under reordering of 0-time interactions. Indeed, if no player can distinguish two histories h and h', then the delay transitions must be the same in the two histories. Further, since all players have access to global time it must be the same actions that are performed in h and h' in every 0-duration sub-history Now, the only way in which h and h' can differ is in the interleaving of these 0-duration subhistories.

Lemma 1 *Let $\sigma = (\sigma_1, ..., \sigma_n)$ be a pure strategy profile and let δ, δ' be two schedulers. Further, let $H = \boldsymbol{s_0 c_0} \xrightarrow{a_1, I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k, I_k} \boldsymbol{s_k}$ be an interval-timed history such that $a_i \in Act$ for all $1 \leq i < k$ and $a_k \in Plr$. Then*

$$\Pr^{\sigma,\delta}_{\boldsymbol{s_0}}(\mathrm{Cyl}^{\mathrm{ia}}(H)) = \Pr^{\sigma,\delta'}_{\boldsymbol{s_0}}(\mathrm{Cyl}^{\mathrm{ia}}(H))$$

Further, this probability is either 0 or 1.

Proof. First note that before a delay transition happens there can be no loops in any of the local histories $\pi_j(H)$. When we restrict to histories up to the first delay transition, a pure strategy for player j can be considered to be simply a maximal sequence $(a^j_1, s^j_1)...(a^j_{\ell_j}, s^j_{\ell_j})$ of choices of player j. This is because for any strategy σ_j of player j, the j-play moves along the unique path in the module of player j that is chosen by player j at each step. The only thing player j can observe is whether the transition (that he chose) is taken or not.

Assume now that we have a pure strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$. This induces such a sequence of choices for each player. The longest possible j-play that can occur under σ before the first delay transition is $\rho^j = s_0(j)(a^j_1, s^j_1), \xrightarrow{a^j_1, 0} s^j_1 ... \xrightarrow{a^j_{\ell_j}, 0} s^j_{\ell_j}$. Further, all possible j-plays are prefixes of ρ^j. We now suppose for contradiction that there are two histories h, h' consistent with σ ending after the first delay transition such that there exists a player j with $|\pi_j(h)| \neq |\pi_j(h')|$. Since $\pi_j(h)$ and $\pi_j(h')$ are both prefixes of ρ^j either $\pi_j(h)$ is a proper prefix of $\pi_j(h')$ or the other way around. Without loss of generality let $u = |\pi_j(h)| < |\pi_j(h')|$.

Suppose $\sigma_j(\pi_j(h)) = (b_0, s)$. Now, player j has not been able to synchronize on $b_0 \in Act$ in the last state of h. However, he has been able to synchronize on it along h' due to different choices of the scheduler. Note that at earlier points on h he might have synchronized on b_0 a number of times already. Let this number of times be $c_0 \in \mathbb{N}$. This means that all players capable of synchronizing on b_0 must

have done so exactly c_0 times along h. At least one of these players, let us call him j_1, must have stopped before committing to synchronize on b_0 the (c_0+1)th time, because otherwise the play would have progressed since no action can be enabled when a delay transition takes place. Note that each of these players are willing to perform b_0 at least c_0+1 times at some point since this happens in h'. Let j_1 be committed to synchronize on action $b_1 \neq b_0$ when the play stops in h and suppose he has already synchronized c_1 times on b_1 before this point. We can now perform the same reasoning again to find a player j_2 that has stopped in h before reaching the point where he is ready to synchronize on b_1 for the (c_1+1)th time. At this point he is committed to performing the action b_2 for the (c_2+1)th time. This reasoning gives us an infinite sequence j_i of players committed to actions b_i in the last state after having performed the action b_i exactly c_i times before.

We now introduce a partial order $\preceq \subseteq (Act \times \mathbb{N})^2$ on elements $(a,d) \in (Act \times \mathbb{N})$ such that there exists j so action a occurs at least d times in h'. The relation is defined such that $(a,d) \preceq (a',d')$ if there is a player j such that a occurs d times on h' before a' occurs d' times on h' (and a' actually does occur d' times at some point on h'). It is reflexive and transitive because all players that has an action in their alphabet must commit in order to synchronize on it. Anti-symmetry follows from this and the fact that ρ^j is linear for every j.

Now, we have that $(b_{i+1}, c_{i+1}) \prec (b_i, c_i)$ for all $i \geq 0$. This is the case firstly because $\pi_{j_i}(h)$ is a prefix of $\pi_{j_i}(h')$ for the players giving rise to the sequence of b_i's. Secondly, because $b_{i+1} \neq b_i$. This means that

$$(b_0, c_0) \succ (b_1, c_1) \succ (b_2, c_2) ...$$

is an infinite strictly decreasing sequence. Since $\succ$ is only defined on a finite number of elements this gives a contradiction. Thus, $|\pi_j(h)| = |\pi_j(h')|$. Further, $\pi_j(h) = \pi_j(h')$ since one is a prefix of the other. Since j was chosen arbitrarily, this means that the local history $\pi_j(h)$ before the delay transition cannot be changed by any scheduler and is uniquely determined by the pure strategies. From this, the lemma follows. □

We now extend to arbitrary events in $\mathcal{I}$ by applying the result above. However, first we need some notation. If $H = \boldsymbol{s_0 c_0} \xrightarrow{a_1, I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k, I_k} \boldsymbol{s_k}$ is an interval-timed history then the set of histories $h \in [H]$ with specific delays $d_1, ..., d_v$ on the delay transitions is denoted $[H]^{d_1,...,d_v}$. Note that this set is finite.

Lemma 2 *Let $E \in \mathcal{I}$, $\boldsymbol{s_0} \in S$, $\sigma = (\sigma_1, ..., \sigma_n)$ be a pure strategy profile and δ, δ' be two schedulers. Then*

$$\Pr^{\sigma,\delta}_{\boldsymbol{s_0}}(E) = \Pr^{\sigma,\delta'}_{\boldsymbol{s_0}}(E)$$

Proof. We show this by showing that for every time-abstract cylinder $\mathrm{Cyl}^{\mathrm{ia}}(H)$ for an interval-timed history $H = \boldsymbol{s_0 c_0} \xrightarrow{a_1, I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k, I_k} \boldsymbol{s_k}$ such that

$a_k \in Plr$ if $k > 0$, every pure strategy profile σ and every pair δ, δ' of schedulers we have

$$\Pr_{\boldsymbol{s_0}}^{\sigma,\delta}(\mathrm{Cyl}^{\mathrm{ia}}(H)) = \Pr_{\boldsymbol{s_0}}^{\sigma,\delta'}(\mathrm{Cyl}^{\mathrm{ia}}(H))$$

First, for any scheduler δ we have

$$\Pr_{\boldsymbol{s_0}}^{\sigma,\delta}(\mathrm{Cyl}^{\mathrm{ia}}(H)) = \sum_{H' \sim H} \Pr_{\boldsymbol{s_0}}^{\sigma,\delta}(\mathrm{Cyl}(H'))$$

since these cylinder sets are disjoint.

Suppose that $a_i \in Plr$ for the indices $\alpha_1, ..., \alpha_v$ and $a_i \in Act$ for the indices $\beta_1, ..., \beta_u$. For $1 \leq i \leq v$, let $\ell_i = \inf(I_{\alpha_i})$ and $u_i = \sup(I_{\alpha_i})$. Next, consider fixed delays $d_1, ..., d_v$ and let $[H]^{d_1,...,d_v} = \{h^1, ..., h^r\}$. For $1 \leq m \leq r$ we denote

$$h^m = \boldsymbol{s_0^m c_0^m} \xrightarrow{a_1^m, t_1} \boldsymbol{s_1^m c_1^m} \cdots \xrightarrow{a_k^m, t_k} \boldsymbol{s_k^m}$$

where the timestamps t_i are induced by the delays $d_1, ..., d_v$ on the delay transitions. Now, since every interval-timed history $H' \sim H$ contains the same intervals and same delay transitions we have

$$\begin{aligned}
&\Pr_{\boldsymbol{s_0}}^{\sigma,\delta}(\mathrm{Cyl}^{\mathrm{ia}}(H)) \\
&= \sum_{H' \sim H} \Pr_{\boldsymbol{s_0}}^{\sigma,\delta}(\mathrm{Cyl}(H')) \\
&= \int_{\ell_1}^{u_1} ... \int_{\ell_v}^{u_v} \prod_{1 \leq i \leq v} Q(s_{\alpha_i}^m, s_{\alpha_i+1}^m) \cdot e^{-E(d_i)} \\
&\cdot \sum_{h^m \in [H]^{d_1,...,d_v}} \Bigg(\prod_{1 \leq i \leq u} \delta(h_{\leq \beta_i - 1}^m, c_{\beta_i - 1}^m)(a_{\beta_i}^m) \\
&\cdot \prod_{0 \leq i < k} \prod_{j \in \mathrm{Sync}(a_i^m)} \sigma_j(\pi_j(h_{\leq i}^m))(c_i^m(j)) \Bigg) \, \mathrm{d}d_v \cdots \mathrm{d}d_1
\end{aligned}$$

We will now show by induction on v that for any fixed delays $d_1, ..., d_v$

$$\begin{aligned}
p^{d_1,...,d_v} = &\sum_{h^m \in [H]^{d_1,...,d_v}} \Bigg(\prod_{1 \leq i \leq u} \delta(h_{\leq \beta_i - 1}^m, c_{\beta_i - 1}^m)(a_{\beta_i}^m) \\
&\prod_{0 \leq i < k} \prod_{j \in \mathrm{Sync}(a_i^m)} \sigma_j(\pi_j(h_{\leq i}^m))(c_i^m(j)) \Bigg)
\end{aligned}$$

either equals 0 or 1 independently of the scheduler δ. This implies that $\Pr_{\boldsymbol{s_0}}^{\sigma,\delta}(\mathrm{Cyl}^{\mathrm{ia}}(H))$ is independent of the scheduler and thus proves the lemma.

For the base case suppose that $v = 0$. Then the only possibility is $H = \boldsymbol{s_0}$ since H cannot end with an action transition. In this case it is immediate that $p^{d_1,...,d_v} = 1$.

For the inductive case suppose that $v > 0$ and that it holds for all H' with less than v delay transitions. Note that $s_{\alpha_{v-1}}$ is the same state for every $h \in$

$[H]^{d_1,\ldots,d_v}$. Thus, we have $[H]^{d_1,\ldots,d_v} = \{h \cdot h' \mid h \in [H_{\leq \alpha_{v-1}}]^{d_1,\ldots,d_{v-1}}$ and $h' \in [H_{\geq \alpha_{v-1}}]^{d_v}\}$. That is, the set of histories in $[H]^{d_1,\ldots,d_v}$ is obtained by gluing together every prefix up to $s_{\alpha_{v-1}}$ with every suffix starting in $s_{\alpha_{v-1}}$. For two histories h^m and h^n in these two sets we denote their concatenation (where final state of h^m is merged with initial state of h^n) by h^{mn}. This gives us

$$
\begin{aligned}
p^{d_1,\ldots,d_v} = & \sum_{h^m \in [H_{\leq \alpha_{v-1}}]^{d_1,\ldots,d_{v-1}}} \sum_{h^n \in [H_{\geq \alpha_{v-1}}]^{d_v}} \Bigg(\\
& \prod_{i:\beta_i < \alpha_{v-1}} \delta(h^m_{\leq \beta_i - 1}, c^m_{\beta_i - 1})(a^m_{\beta_i}) \\
& \prod_{0 \leq i < \alpha_{v-1}} \prod_{j \in \mathrm{Sync}(a^m_i)} \sigma_j(\pi_j(h^m_{\leq i}))(c^m_i(j)) \\
& \prod_{i:\beta_i > \alpha_{v-1}} \delta(h^{mn}_{\leq \beta_i - 1}, c^{mn}_{\beta_i - 1})(a^{mn}_{\beta_i}) \\
& \prod_{\alpha_{v-1} \leq i < \alpha_v} \prod_{j \in \mathrm{Sync}(a^{mn}_i)} \sigma_j(\pi_j(h^{mn}_{\leq i}))(c^{mn}_i(j)) \Bigg) \\
= & \sum_{h^m \in [H_{\leq \alpha_{v-1}}]^{d_1,\ldots,d_{v-1}}} \\
& \cdot \prod_{i:\beta_i < \alpha_{v-1}} \delta(h^m_{\leq \beta_i - 1}, c^m_{\beta_i - 1})(a^m_{\beta_i}) \\
& \cdot \prod_{0 \leq i < \alpha_{v-1}} \prod_{j \in \mathrm{Sync}(a^m_i)} \sigma_j(\pi_j(h^m_{\leq i}))(c^m_i(j)) \\
& \cdot \sum_{h^n \in [H_{\geq \alpha_{v-1}}]^{d_v}} \Bigg(\prod_{i:\beta_i > \alpha_{v-1}} \delta(h^{mn}_{\leq \beta_i - 1}, c^{mn}_{\beta_i - 1})(a^{mn}_{\beta_i}) \\
& \prod_{\alpha_{v-1} \leq i < \alpha_v} \prod_{j \in \mathrm{Sync}(a^{mn}_i)} \sigma_j(\pi_j(h^{mn}_{\leq i}))(c^{mn}_i(j)) \Bigg)
\end{aligned}
$$

If we can show that the second sum is either 0 or 1 independently of the scheduler then we can apply the induction hypothesis on the remaining part. Note that no player can distinguish between the prefixes h^m since they are in the same equivalence class. Thus, using pure strategies the players can only base their decision on what happens after reaching $s^{mn}_{\alpha_{v-1}}$. Now, using the same technique as in the proof of Lemma 1 the result follows. □

From Lemma 2 we know that the probabilities of events in $\mathcal{I}$ are independent of the scheduler when the strategy profile is pure. Using this we can show that this is also the case for non-pure strategy profiles. The idea of the proof is similar to the proof of Kuhn's Theorem [3] behavioural strategies are shown equivalent in perfect recall extensive-form games. This intuition is applied in the 0-duration subhistories.

[3] H. W. Kuhn, "Extensive games and the problem of information," Annals of Mathematics Studies, vol. 28, 1953

Theorem 1. Let $E \in \mathcal{I}$, let $\boldsymbol{s_0} \in S$ be a state, let $\sigma = (\sigma_1, ..., \sigma_n)$ be a strategy profile and δ, δ' be two schedulers. Then

$$\Pr^{\sigma,\delta}_{\boldsymbol{s_0}}(E) = \Pr^{\sigma,\delta'}_{\boldsymbol{s_0}}(E)$$

Proof. We show this by showing that for every time-abstract cylinder $\mathrm{Cyl}^{\mathrm{ia}}(H)$ for an interval-timed history $H = \boldsymbol{s_0 c_0} \xrightarrow{a_1, I_1} \boldsymbol{s_1 c_1} \cdots \xrightarrow{a_k, I_k} \boldsymbol{s_k}$ such that $a_k \in \mathit{Plr}$ if $k > 0$, every strategy profile σ and every pair δ, δ' of schedulers we have

$$\Pr^{\sigma,\delta}_{\boldsymbol{s_0}}(\mathrm{Cyl}^{\mathrm{ia}}(H)) = \Pr^{\sigma,\delta'}_{\boldsymbol{s_0}}(\mathrm{Cyl}^{\mathrm{ia}}(H))$$

We use the same notation as in the proof of Lemma 2. Again we have

$$\begin{aligned}
&\Pr^{\sigma,\delta}_{\boldsymbol{s_0}}(\mathrm{Cyl}^{\mathrm{ia}}(H)) \\
&= \int_{\ell_1}^{u_1} ... \int_{\ell_v}^{u_v} \prod_{1 \le i \le v} Q(s^m_{\alpha_i}, s^m_{\alpha_i+1}) \cdot e^{-E(d_i)} \\
&\cdot \sum_{h^m \in [H]^{d_1,\ldots,d_v}} \Bigg(\prod_{1 \le i \le u} \delta(h^m_{\le \beta_i - 1}, c^m_{\beta_i - 1})(a^m_{\beta_i}) \\
&\cdot \prod_{0 \le i < k} \prod_{j \in \mathrm{Sync}(a^m_i)} \sigma_j(\pi_j(h^m_{\le i}))(c^m_i(j)) \Bigg) \, \mathrm{d}d_v \cdots \mathrm{d}d_1
\end{aligned}$$

For fixed delays $d_1, ..., d_v$ we will show that a discrete probability distribution over a finite set of pure strategies gives rise to the same probability as above. As this is independent of the scheduler by Lemma 2 so is the probability for the mixed strategies. As this holds for all delays, the Theorem follows.

Now, for each history $h^m \in [H]^{d_1,\ldots,d_v}$ we define a pure strategy profile σ^{h^m} that plays according to h^m as well as a probability p^{h^m} defined by

$$p^{h^m} = \prod_{0 \le i < k} \prod_{j \in \mathrm{Sync}(a^m_i)} \sigma_j(\pi_j(h^m))(c^m_i(j))$$

If $p = \sum_{h^m \in [H]^{d_1,\ldots,d_v}} p^{h^m} < 1$ then define a pure strategy σ'' that plays such that a $h \in [H]^{d_1,\ldots,d_v}$ is not possible. Further, define $p^{\sigma''} = 1 - p$. Now, consider the experiment of using scheduler δ and picking either one of the strategy profiles σ^{h^m} with probability p^{h^m} or σ'' with probability $p^{\sigma''}$ and applying these strategies. The probability that a prefix of the play is in $[H]^{d_1,\ldots,d_v}$ in this experiment is independent of the scheduler because of Lemma 2. We will show that it is in fact equal to $\Pr^{\sigma,\delta}_{\boldsymbol{s_0}}(\mathrm{Cyl}^{\mathrm{ia}}(H))$. Indeed, the probability is

$$\sum_{\sigma^{h^m}} p^{h^m} \sum_{h^m \in [H]^{d_1,\ldots,d_v}} \cdot \Bigg(\prod_{1 \le i \le v} Q(s^m_{\alpha_i}, s^m_{\alpha_i+1}) \cdot e^{-E(d_i)}$$

$$\cdot \prod_{1\le i\le u} \delta(h^m_{\le\beta_i-1}, c^m_{\beta_i-1})(a^m_{\beta_i})$$
$$\cdot \prod_{0\le i<k} \prod_{j\in \mathrm{Sync}(a^m_i)} \sigma_j^{h^m}(\pi_j(h^m_{\le i}))(c^m_i(j))\Big)$$
$$= \prod_{1\le i\le v} Q(s^m_{\alpha_i}, s^m_{\alpha_i+1}) \cdot e^{-E(d_i)}$$
$$\cdot \sum_{h^m\in[H]^{d_1,\dots,d_v}} \Big(p^{h^m} \cdot \prod_{1\le i\le u} \delta(h^m_{\le\beta_i-1}, c^m_{\beta_i-1})(a^m_{\beta_i})\Big)$$

By inserting the expression for p^{h^m} the result follows. □

C Full proof of Proposition 1

Proposition 1. For a DEC-POMDP $\mathcal{P}$ with n players and target set T we can construct in polynomial time a distributed IMC $\mathcal{G}$ with $n+1$ players and target set T' such that

$$\exists \boldsymbol{\sigma} : \Pr^{\boldsymbol{\sigma}}_{\mathcal{G}}(\diamond T) = p \iff \exists \boldsymbol{\sigma'} : \Pr^{\boldsymbol{\sigma}'}_{\mathcal{P}}(\diamond T') = p.$$

Proof. Let us fix a DEC-POMDP $\mathcal{P} = (S, Plr, (Act_i)_{1\le i\le n}, \delta, (\mathcal{O}_i)_{1\le i\le n}, O, s^{in})$ with n players $Plr = \{1, ..., n\}$. Further, let $Act_i = \{a_{i1}, ..., a_{im_i}\}$ and $\mathcal{O}_i = \{o_{i1}, ..., o_{i\ell_i}\}$ for player $i \in Plr$. The distributed IMC $\mathcal{G}$ has $n+1$ modules, one module for each player in $\mathcal{P}$ and the *main* module responsible for their synchronization. Intuitively,

- the module of every player i stores the last local observation in its state space. Every step of $\mathcal{P}$ is modelled as follows: The player *outputs* to the main module the action it chooses and then *inputs* from the main module the next observation.
- The main module stores the global state in its state space. Every step of $\mathcal{P}$ corresponds to the following: The main module *inputs* the actions of all players one by one, then it randomly picks the new state and new observations according to the rules of $\mathcal{P}$ based on the actions collected. The observations are lastly *output* to all players, again one by one. The main module is constructed so in every state there is at most one action transition. Thus, there is only one trivial strategy that cannot influence anything.

The construction of module for player $i \in \{1, ..., n\}$ is illustrated in Figure 4 along with constructions for input and output. The interesting part is the inputting mechanism. Inputting an action from the set $\{a_1, \dots, a_r\}$ when in a state s is done as follows. Instead of waiting in s, the player travels by delay transitions in a round-robin fashion through a cycle of r states, where in the i-th state, only the action a_i is available. This way, the player cannot influence anything and must input the action that comes with probability 1.

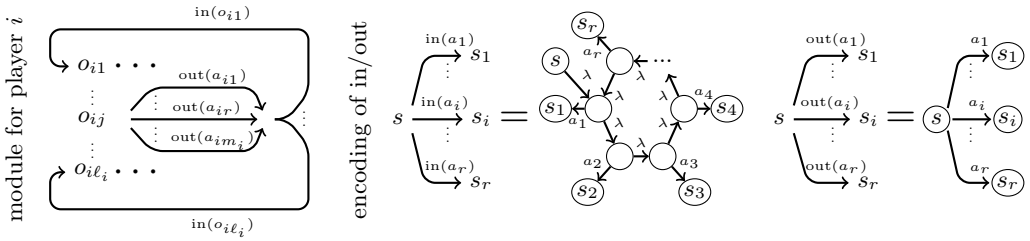

Fig. 4: Module for player i on the left. Input and output encoding to the right.

In the same fashion, the main module is constructed such that the extra player $n+1$ has at most one possible choice in each state and thus, this player controls nothing except what he is forced to do by the structure of his module. Thus, he has no choice but to enforce the rules of $\mathcal{P}$ in $\mathcal{G}$. In the main module the current state s of the DEC-POMDP $\mathcal{P}$ is remembered at all times. Player $n+1$ then goes through n steps inputting an action for each player. These are saved in the state of the module as well. When all actions have been collected, a random choice according to δ is made. This determines the successor state s'. Afterwards, a random choice of observations for the different players is done according to O. Then, these observations are outputted to the other players again in a sequence of n steps. The play proceeds for an infinite number of rounds, thereby modelling the DEC-POMDP.

The initial global state in $\mathcal{G}$ is known to all players and each initial local state is left one by one as the players output their first action. In addition, $\boldsymbol{s_0}(n+1) = s^{in}$ as the main module mimics $\mathcal{P}$. The target set T' in $\mathcal{G}$ is given by all global states such that the main module is in a state corresponding to a state in T.

Since none of the players can affect the timing in the distributed IMC created (they always need to choose actions immediately and are never allowed to redecide) none of them can gain any information that was not already there in the DEC-POMDP. Therefore, there exists a strategy profile $\boldsymbol{\sigma}$ in $\mathcal{P}$ with $\Pr^{\boldsymbol{\sigma}}_{\mathcal{P}}(\diamond T) = p$ if and only if there exists a strategy profile $\boldsymbol{\sigma'}$ in $\mathcal{G}$ with $\Pr^{\boldsymbol{\sigma}'}_{\mathcal{G}}(\diamond T') = p$. Note that this is also the case since there is never more than one enabled action at a time in the main module. Thus, player $n+1$ does not have any influence. □

D Full proof of Theorem 3

Proposition 3. The qualitative existence problem for DEC-POMDPs with $n \geq 2$ players is undecidable.

Proof. We do a reduction from the non-halting problem of a deterministic Turing machine $M = (Q, q_0, \Sigma, \Delta, B, F)$ that never writes the blank symbol B and starts with a blank input tape. Further, it never moves to the left of the initial

state. Here, Q is the finite set of control states, q_0 is the initial state, Σ is the tape alphabet, $\Delta : Q \times \Sigma \to Q \times \Sigma \times \{L, R\}$ is the transition function and F is the set of accepting states. From M we construct a DEC-POMDP $\mathcal{P} = (S, \{1, 2\}, (Act_1, Act_2), \delta, (\mathcal{O}_1, \mathcal{O}_2), O)$ with two players $Plr = \{1, 2\}$ such that M does not halt if and only if player 1 and 2 have strategies σ_1 and σ_2 such that the probability of reaching a target set $T \subseteq S$ is 1. The overall structure of $\mathcal{P}$ can be seen in Figure 5 but without the details of the sub-module $\mathcal{P}'$.

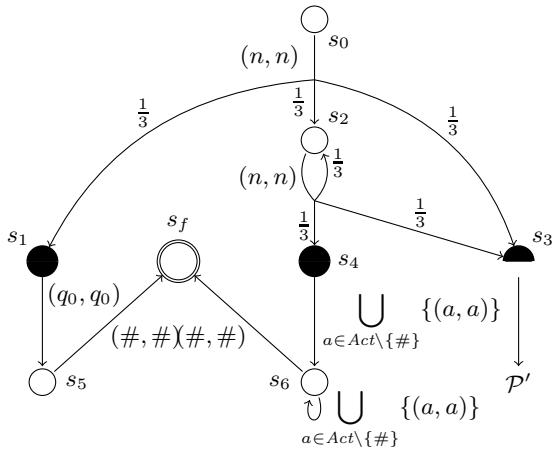

Fig. 5: Overall structure of $\mathcal{P}$ without the details of $\mathcal{P}'$.

In this game, both players have two possible observations. In all states except s_1, s_3, s_4 player 1 receives one observation (illustrated with top-half of states filled with black) and in all other states he will receive another observation (top-half of states being white). Player 2 gets one observation for states s_1, s_4 as well as in some states in $\mathcal{P}'$ (when bottom-half is black). In all other states he receives another observation.

The play starts in s_0 and with probability 1 either each player will receive the black observation exactly once during the game or the play will go to $s_f \in F$ from s_2 at some point. If the play goes to s_1 or s_4 the players will receive the observation at the same time and if the play goes to s_3 then player 2 will receive the observation in the next state. We will show that in order to have strategies σ_1, σ_2 that can make sure to reach the target set with probability 1 two things must be satisfied

1. σ_i must prescribe playing the jth configuration of the Turing machine M (in a format described below) if player $i \in \{1, 2\}$ receives the black observation in the jth step after the play leaves s_0
2. M does not halt

Further, when these two properties are satisfied, players 1 and 2 can indeed reach the target set with probability 1 by applying σ_1 and σ_2. We now explain

what we mean by playing a configuration of M. Let $(q, w, j) \in Q \times \Sigma^* \times \mathbb{N}$ be a configuration of M where q is the current control state, w is the current non-blank part of the tape contents and j is the current position of the tape head. By playing configuration (q, w, j) we mean playing the following sequence of actions (where possible actions are corresponding to control states, tape symbols and # as an end of non-blank tape marker)

$$w_1 w_2 ... w_{j-1} q w_j w_{j+1} ... w_{|w|} \#$$

In other words, the contents of the tape is played one symbol at a time and the control state is prior to contents of the tape cell that the tape head points to.

The DEC-POMDP $\mathcal{P}$ is constructed such that

1. If the play goes to s_1 both players must play the first configuration of M
2. If the play goes to s_4 both players must play the same sequence of symbols and ending with a #
3. If the play goes to s_3 the two players must play configurations C_1 and C_2 respectively such that C_2 is a successor configuration of C_1 according to the transition function of M. Further, this must be done with an offset of 1 step because player 2 receives the black observation one step later than player 1. Finally, the play will go to a sink state if they play a halting state as part of the configurations, making it impossible to reach T afterwards. This can be done with a finite module $\mathcal{P}'$ [5].

By induction we can show that if player 1 and 2 have strategies σ_1 and σ_2 to ensure reaching T with probability 1 then they must play the jth configuration of M if receiving the black observation in the jth step after the play leaves s_0. Indeed, if one of the players gets the black observation for $j = 1$ then s_1 is a possible state. Thus, both players must play the first configuration when $j = 1$. Now, suppose as induction hypothesis it is true for some j. Now, if player 2 receives the black observation in the $(j + 1)$th step then the play might have passed through s_3 in which case player 1 will have received the black observation in the jth step. Thus, to make sure to reach T in this case player 2 must player a successor configuration of what player 1 plays. Thus, player 2 must play the $(j + 1)$th configuration of M when getting the black observation in the $(j + 1)$th step. Now, if player 1 gets the black observation after $j + 1$ steps, then it is possible that the play is in s_4. If this is the case then player 2 will also have received the black observation in the $(j + 1)$th step. Since player 2 will then play the $(j + 1)$th configuration of M then player 1 has to do this as well in order to reach T since the players must play the same sequence of symbols to reach T from s_4. This concludes the induction step.

We have now shown that if player 1 and 2 have strategies σ_1 and σ_2 to ensure reaching T with probability 1 then they must play the jth configuration of M when receiving the black observation in the jth step after the play leaves s_0. We now need to show that if M does not halt, then applying these strategies will actually ensure reaching T with probability 1 and if M does halt, then applying

these strategies will not ensure reaching T with probability 1 (and thus, in this case, no strategies can ensure this).

In the case where M does not halt, the play will reach s_f, s_1, s_3 or s_4 with probability 1. Suppose it reaches s_1, s_3 or s_4 at step j. Then both players play the appropriate configuration and reach T with probability 1 since none of them will play a halting state at any point.

Suppose on the other hand that M halts after j steps. Now the play will reach s_4 after j steps with positive probability. And since the players using σ_1 and σ_2 play the jth configuration of M when this happens, they will play a halting state with positive probability. Thus, they cannot ensure reaching T with probability 1 in this case.

In total this means the players have strategies to reach T with probability 1 if and only if M does not halt. □

E Proof of Proposition 2

Proposition 2. The value of $\mathcal{G}$ equals to the value of $\mathcal{M}_\mathcal{G}$, i.e.

$$\sup_{\boldsymbol{\sigma}} \Pr^{\boldsymbol{\sigma}}(\diamond T) \;=\; \sup_{\pi} \Pr^{\pi}(\diamond T).$$

The proof goes by three technical steps. First we define a stronger class of *synchronization* strategies that can observe the whole global state whenever in a synchronization state that have the same power as the strategies in $\mathcal{M}_\mathcal{G}$. Second, we show that any strategy profile can be mimicked by a synchronization strategy. The crucial step is the third. We show how any synchronization strategy can be (up to an arbitrary error) emulated by standard strategies in a distributed IMC. These strategies obtain (with high probability) full-observation by delaying.

E.1 Synchronization strategies

A *global strategy* is a measurable function $\theta : \mathit{Hist} \to \Delta(\mathcal{C})$ that takes new choices only when allowed, i.e. that for any global history h of the form $h = h' \cdot \boldsymbol{c} \xrightarrow{a,t} \boldsymbol{s}$ and any global choice $\boldsymbol{c'}$ such that $\theta(h)(\boldsymbol{c'}) > 0$ we have $\boldsymbol{c}(j) = \boldsymbol{c'}(j)$ if $j \notin \mathrm{Sync}(a)$. For this section, we call standard strategies *local* to stress the difference.

Furthermore, we call a global strategy a *synchronization* strategy if it, intuitively, disregards non-local knowledge after the last synchronization state. Formally, the following condition needs to be satisfied for any pair of global histories h and h' and any player $j \in \{1, 2\}$. Let $h = h_1 \cdot s \cdot h_2$ and $h' = h'_1 \cdot s' \cdot h'_2$ where h_2 and h'_2 are the sequences after the last visit of a synchronization state (s and s'). If (1) $h_1 \cdot s = h'_1 \cdot s'$, and (2) $\pi_j(h_2) = \pi_j(h_2)$, we have $\theta(h)(c) = \theta(h')(c)$ for any $c \in \mathcal{C}_j$ where $\theta(h)(c) := \sum_{\boldsymbol{c}, \boldsymbol{c}(j)=c} \theta(h)(\boldsymbol{c})$.

Note that a synchronization strategy replaces a complete profile of local strategies. The definition of the probability measure $\Pr^{\theta,\delta}_{\boldsymbol{s_0}}$ induced by a synchronization strategy θ and a scheduler δ goes along the same lines as for a strategy

profile. The only difference is that we replace the term St_i by the much simpler term $St'_i = \theta(h_i)(\boldsymbol{c_i})$. By observing the whole synchronization state the strategy θ can completely avoid that a scheduler δ has any power. Indeed, the strategy θ can make sure that $\mathrm{En}(\boldsymbol{s})$ always contains at most one action. Similarly to local strategies, we thus simplify the notation by denoting the probability measure $\Pr^{\theta}_{\boldsymbol{s_0}}$.

Lemma 1. *The value of $\mathcal{G}$ w.r.t. synchronization strategies equals to the value of $\mathcal{M}_{\mathcal{G}}$:*

$$\sup_{\pi} \Pr^{\pi}(\diamond T) \;=\; \sup_{\theta} \Pr^{\theta}(\diamond T)$$

Proof. Ad $\leq$: We show that an optimal pure memoryless strategy π induces a synchronization pure memoryless strategy θ with the same value. Let h be a history in $\mathcal{G}$. This gives us a history $\bar{h}$ in the MDP (by removing choices, actions, time, and all private states). Further, let $(\boldsymbol{c}, \sigma_1, \sigma_2) = \pi(\bar{h})$ be the decision of π.

- If $\mathrm{last}(h)$ is a synchronization state, θ takes the choice $\boldsymbol{c}$ (maybe step by step as the choice of each player can be changes only after the player moves).
- If $\mathrm{last}(h)$ is a private state and the player $j \in \{1, 2\}$ needs to take an action, we distinguish two cases:
  - if $\mathrm{last}(h)(j)$ is a private state, player j takes action according to the pure memoryless strategy σ_j;
  - if $\mathrm{last}(h)(j)$ is a synchronization state, player j chooses $\varnothing_j$.

The equality $\Pr^{\pi}(\diamond T) = \Pr\theta(\diamond T)$ can be easily shown by induction in the number of visits to synchronization states. The crucial fact is that the probability to reach in $\mathcal{G}$ by θ a state $(s'_1, s'_2) \in S'$ from a state $(s_1, s_2) \in S'$ only via $S \setminus S'$ coincides with $P((s_1, s_2), \pi((s_1, s_2)))(s'_1, s'_2)$.

Ad $\geq$: For a global history $h \cdot (s_1, s_2)$, we denote by $\Pr^{\theta}_{h \cdot (s_1, s_2)}(\diamond T)$ the probability to reach T with θ after history $h \cdot (s_1, s_2)$ has already passed. We show the statement in three steps.

First we argue that

$$\sup_{\theta} \Pr^{\theta}_{h \cdot (s_1, s_2)}(\diamond T) = \sup_{\theta} \Pr^{\theta}_{h' \cdot (s_1, s_2)}(\diamond T)$$

for any histories $h \cdot (s_1, s_2)$ and $h' \cdot (s_1, s_2)$ ending in the same synchronization state. Let us assume that for some θ' we have that $\Pr^{\theta'}_{h \cdot (s_1, s_2)}(\diamond T)$ is strictly greater than the value for $h'(s_1, s_2)$. Since only the last synchronization state has impact on the further evolution, we could define a strategy θ'' to play in $h' \cdot (s_1, s_2)$ in the same way as θ' plays in $h \cdot (s_1, s_2)$. This way, we get the same reachability probability yielding a contradiction.

Second, we can thus define the *value* of a synchronization state (s_1, s_2) simply as

$$val(s_1, s_2) = \sup_{\theta} \Pr^{\theta}_{(s_1, s_2)}(\diamond T)$$

and we can easily obtain equations

$$val(s_1, s_2) = \sup_{\theta} \sum_{(s_1', s_2') \in S'} \Pr^{\theta}_{(s_1,s_2)}(R(s_1', s_2')) \cdot val(s_1', s_2').$$

where $R(s_1', s_2')$ are the plays that reach (s_1', s_2') via states in $S \setminus S'$. Because a synchronization strategy needs to play in each such fragment exactly as a profile of local strategies, it is in turn equal to

$$= \sup_{(\sigma_1,\sigma_2)} \sum_{(s_1', s_2') \in S'} \Pr^{(\sigma_1,\sigma_2)}_{(s_1,s_2)}(R(s_1', s_2')) \cdot val(s_1', s_2')$$

which can be decomposed to two independent one-player games

$$= \sup_{\sigma_1,\sigma_2} \sum_{(s_1', s_2') \in S'} \Pr^{\sigma_1}_{s_1}(R(s_1')) \cdot \Pr^{\sigma_2}_{s_2}(R(s_2')) \cdot val(s_1', s_2').$$

Each such one player game is equivalent to a continuous-time Markov decision process where any achievable vector of (time-unbounded) reachability probabilities of terminal states can be equivalently achieved in the embedded discrete-time Markov chain. The set of achievable vectors of reachability probabilities forms a polytope with the corners given by pure memoryless strategies. One can easily see that the supremum above is always obtained for some corners of the two polytopes, i.e. for pure memoryless strategies. □

E.2 Proof of $\leq$:

We show that for every profile $\boldsymbol{\sigma}$ of local strategies there is a synchronization strategy θ such that

$$\sup_{\boldsymbol{\sigma}} \Pr^{\boldsymbol{\sigma}}(\diamond T) \leq \Pr^{\theta}(\diamond T)$$

The proof follows from the fact that for a profile $\boldsymbol{\sigma}$ of local strategies there is a synchronization strategy θ inducing equivalent probability measures. Indeed, for decisions of every player i the synchronization strategy disregards the global knowledge and decides only based on the ith local projection of the history exactly as σ_i.

E.3 Proof of $\geq$:

Finally, we show that using the optimal strategy π', we can define a sequence of local strategy profiles $\boldsymbol{\sigma}^1, \boldsymbol{\sigma}^2, \ldots$ such that

$$\lim_{i \to \infty} \Pr^{\boldsymbol{\sigma}^i}(\diamond T) = \Pr^{\pi'}(\diamond T).$$

Let us fix $i \in \mathbb{N}$. In order to define the strategy profile $\boldsymbol{\sigma}^i$, we need some auxiliary notions. We assume w.l.o.g. that $S_1' = \{0, \ldots, |S_1| - 1\}$ and $S_2' =$

$\{0, \ldots, |S_2| - 1\}$. Let λ be the minimal rate of the delay transition over all synchronization states of both players and K_i be the minimal integer such that $e^{-\lambda K_i} \leq 1/(4i \cdot (|S_1'| + |S_2'|))$. We split for each count of synchronization steps $k \in \mathbb{N}_0$ the time into *phases* of length $|S_1'| \cdot 2|S_2'| \cdot (K_i + k/\lambda)$ each composed of $|S_1'| \cdot 2|S_2'|$ *slots* of length $\ell_{i,k} = (K_i + k/\lambda)$. For any $t \in \mathbb{R}$ we define functions

- $\mathrm{row}_{i,k}(t) := \lfloor t \;/\; 2|S_2'| \cdot \ell_{i,k} \rfloor \mod |S_1'|$,
- $\mathrm{col}_{i,k}(t) := \lfloor t \;/\; \ell_{i,k} \rfloor \mod 2|S_2'|$.

that give the current row and column in a synchronization table depicted in the figure on page 14.
Let h be a history that includes k synchronization actions.

- Let h end in a synchronization state (s_1, s_2) at time t, and let the timed guesses be $\bar{s}_1 = \mathrm{row}_{i,k}(t)$ and $\bar{s}_2 = \lfloor \mathrm{col}_{i,k}(t) \;/\; 2 \rfloor$. Furthermore, let $(c_1^{\bar{s}_1,\bar{s}_2}, c_2^{\bar{s}_1,\bar{s}_2})$ be the optimal choices of π' for players 1 and 2 in state $(\bar{s}_1, \bar{s}_2)$. Finally, we set a proposition sync to true iff $\mathrm{col}_{i,k}(t) \mod 2 = 0$. We define for each player (if the player is allowed to change the choice at the moment)

$$\sigma_1^i(h) = \begin{cases} c_1^{\bar{s}_1,\bar{s}_2} & \text{if } \bar{s}_1 = s_1 \text{ and sync,} \\ \varnothing_1 & \text{otherwise;} \end{cases}$$

$$\sigma_2^i(h) = \begin{cases} c_2^{\bar{s}_1,\bar{s}_2} & \text{if } \bar{s}_2 = s_2 \text{ and sync,} \\ \varnothing_2 & \text{otherwise;} \end{cases}$$

- If h ends in a private state (s_1, s_2), let t be the time of last synchronization. We define again the timed guesses by $\bar{s}_1 = \mathrm{row}_{i,k}(t)$ and $\bar{s}_2 = \lfloor \mathrm{col}_{i,k}(t) \;/\; 2 \rfloor$. Then σ_j^i plays in s_j as the pure memoryless strategy σ_j chosen by the optimal strategy π' in state $(\bar{s}_1, \bar{s}_2)$.

Lemma 7 *For any $i \in \mathbb{N}$ we have*

$$|\Pr_s^{\pi'}(\diamond T) - \Pr_s^{\boldsymbol{\sigma}^i}(\diamond T)| \;\leq\; \frac{1}{i}.$$

Proof. Let us fix $i \in \mathbb{N}$ and $k \in \mathbb{N}_0$. From the definition of $\ell_{i,k}$ we have that the probability $err_{i,k}$ that a choice is *not* updated within the length of one slot is

$$err_{i,k} \leq e^{-\lambda(K_i + k/\lambda)} \leq e^{-\lambda K_i}/2^k \leq 1/(4i \cdot (|S_1| + |S_2|) \cdot 2^k).$$

Thus, the probability $corr_{i,k}^1$ that the correct synchronization is achieved within one phase is at least $corr_{i,k} \geq 1 - 2err_{i,k}$ and the probability $inco_{i,k}^1$ that an incorrect synchronization is achieved within one phase is at most $inco_{i,k}^1 \leq err_{i,k} \cdot (|S_1| + |S_2|)$ (here we bound it by the event that any synchronization choice is not switched off within the next ¬sync slot). With the remaining probability $wait_{i,k}^1 = 1 - corr_{i,k}^1 - inco_{i,k}^1$, no synchronization is achieved within one phase.

The overall probability that the kth synchronization is incorrect is

$$inco_{i,k} = inco^1_{i,k} \cdot \sum_{j=0}^{\infty} \left(wait^1_{i,k}\right)^j \leq inco^1_{i,k} \cdot \sum_{j=0}^{\infty} \left(1 - corr^1_{i,k}\right)^j$$
$$\leq err_{i,k} \cdot (|S_1| + |S_2|) \cdot \sum_{j=0}^{\infty} (2err_{i,k})^j \leq \frac{1}{2 \cdot i \cdot 2^k}.$$

Therefore, the overall probability that *any* synchronization is incorrect is $\sum_{j=0}^{\infty} inco_{i,j} \leq 1/i$. It is easy to see that every profile $\boldsymbol{\sigma}^i$ emulates π', provided every synchronization guarantees that both players have correct knowledge of the other player's state. □